

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE March 1996		3. REPORT TYPE AND DATES COVERED Jan 1/95 - Jan 1/96
4. TITLE AND SUBTITLE An Architecture for Visualization and User Interaction in Parallel Environments			5. FUNDING NUMBERS DAAH04-93-G-0045	
6. AUTHOR(S) E. Mascarenhas and V. Rego				
7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(ES) Department of Computer Sciences Purdue University West Lafayette IN 47907			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSORING / MONITORING AGENCY REPORT NUMBER ARO 33703.1-MA-RIP	
11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.			12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) An application-independent visualization interaction system is proposed with potential for application-binding at any stage of the modeling process. Advantages of this approach include ease of use, flexibility, code reuse, and modularity. Our design ideas are manifest in the DISplay system, a graphical user interaction and display library which can be used with any parallel software. We outline its use in the dynamic display of results from computations in queueing simulations, exemplifying synchronization of multiple remote display requests, and the potential for enhanced parallel simulations. We also describe how it may be used to define customized user interaction dialogs for input from an application, and bi-directional interaction between a user and an application.				
14. SUBJECT TERMS			15. NUMBER OF PAGES	
16. PRICE CODE				
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED		18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED		19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED
20. LIMITATION OF ABSTRACT UL				

19960522 109

An Architecture for Visualization and User Interaction in Parallel Environments *

Edward Mascarenhas
Vernon Rego

Department of Computer Sciences
Purdue University
West Lafayette, IN 47907

Abstract

An application-independent visualization interaction system is proposed, with potential for application-binding at any stage of the modeling process. Advantages of this approach include ease of use, flexibility, code reuse, and modularity. Our design ideas are manifest in the DISplay system, a graphical user interaction and display library which can be used with any parallel software. We outline its use in the dynamic display of results from computations in queueing simulations, exemplifying synchronization of multiple remote display requests, and the potential for enhanced parallel simulations. We also describe how it may be used to define customized user interaction dialogs for input from an application, and bi-directional interaction between a user and an application.

* Research supported in part by NATO-CRG900108, NSF CCR-9102331, ONR-9310233, and ARO-93G0045.

1 Introduction

The methodology of graphical user interaction is well accepted as a useful and constructive modeling aid in simulation applications. In application-specific user interactions with a computation, an interface is typically made part of an application. In contrast to this, we propose the use of an application-independent visualization interaction system with potential for application-binding at any stage of the modeling process. Advantages of this approach include *ease of use* – specialized display knowledge is encapsulated in the system, *flexibility* – the system can be used for a variety of applications, including parallel and distributed simulations, *code reuse* – graphical interaction code need not be reconstructed for different applications, and *modularity* – visualization and interaction portions of the application can be layered on top of the application, facilitating design changes and code modification. Our ideas are manifest in the DISplay system, a graphical user interaction and display library, and associated servers which can be used with any sequential or parallel computation. DISplay was motivated by our studies in Distributed Interactive Simulation on heterogeneous networks [7]. In this paper we focus primarily on the use of DISplay in parallel simulations.

The construction of a user-interface for a given application is a nontrivial task, requiring specialized skills. A good interface enhances the ease of use of a software product, often improving a user's productivity significantly. Further, the success of a software product often depends on the interface it provides to potential users. In developing a user interface for a given application, a designer chooses one of two standard approaches. One approach is to bind the coding of the interface with the coding of the application; the interface becomes an integral part of the application. Though parts of the the interface may be common to other applications, these parts must be recoded when used elsewhere. To get around this, the entire functionality of an interface can be embedded in a high-level library which provides an application developer a tool for rapid interface generation. This is the second approach and is found in tools like Motif [6] and Tk [12]. This tack has several advantages. Application developers require no knowledge of the workings of low levels graphics functions. Distinct applications may link to the same library, using a common interface to make distinct application-specific interfaces. Display and interaction routines contained in the library may be shared by different applications, also known as code-reuse in the software engineering community.

In a typical modeling study, an analyst (user) initially interacts with a simulation to provide input, e.g. by providing parameters or the name of a file containing parameters. As the simulation proceeds, it may be necessary to interact with the simulation and provide input to change its computational trajectory, based on its current trajectory. Such interaction with an executing model may take place at specific or at arbitrary points in the model, depending on the type of interaction required. It is generally useful to an analyst to have a continually updated status report on the progress of an executing model, either via textual displays, graphical displays, or an abstract display representing meaningful progress so that a run may be aborted when a computation goes awry. When a simulation run is done, an analyst usually wants output in the form of tables or graphs. The complexity of such user-interactions is significantly larger in

the parallel simulation, and more generally, parallel computing domains. Simple sequential displays now require synchronization between multiple (distributed) processes, so that textual or graphical results are presented to a user in a consistent manner, obeying causality constraints. Simple user-interface tools for sequential interactions and displays cannot cater to multiple interactions for distributed applications.

1.1 Visual Interactive Simulation and DISplay

Questions concerning the provision of facilities for user interaction with a running simulation, or for the graphical display of simulations results in real-time are not new. Indeed, O'Keefe [10] discusses basic ideas underlying and a methodology for Visual Interactive Simulation (VIS), and Rooks [14] gives a proposal for VIS, outlining general requirements and potentially unifying framework.

Early software systems supporting VIS were generally restricted to animations of application components of simulation models. Later, these systems added user interaction capabilities in various forms. More recent systems supporting VIS include WITNESS [18] — which also allows interactive model building, Cinema [5] and Arena [1] — which are used with the SIMAN simulation language, providing real-time or post-processed animation and building of a model, SIMSCRIPT II.5 [15] — a language which provides an integrated graphical interface called SIMGRAPHICS, and TESS [13, 11] — a system associated with the SLAM simulation language, with facilities for graphical model building, analyzing, graphing and animating model results.

In this paper we describe the DISplay system, which exhibits all of the requirements of VIS to a large degree. Our terminology follows that provided by Rooks [14]. Facilities for *dynamic visualization* of the progress of a simulation in the form of *abstract displays* like histograms and *representative displays* like networks are provided. The result is a system which provides a simple scheme for animating general simulations. DISplay is implemented in C and C++. The library with which the application program is linked is entirely C-based. A programmer may code an application in a variety of languages, provided these can be interfaced with a C library. We have tested the DISplay tool in a number of different environments, including distributed computing environments such as PVM [17] and Conch [19], the *EclIPSe* parallel simulation environment [7], and the process-oriented simulation tool CSIM [16]. Here we describe one example of its use with CSIM. Use of DISplay for the display of real-time performance measures in general *EclIPSe*-based parallel simulations is described in [7]. DISplay has a graphical component that was constructed using OSF/Motif [6], in conjunction with the PEX [4] library for 3D graphs. To use DISplay, all that is required is a workstation hosting an X server, with an optional PEX compatibility (for 3D graphs).

The rest of the paper is organized as follows. Section 2 contains a description of the DISplay system architecture, our design goals and software requirements. In Section 3 we describe design aspects in some detail, also motivating our design decisions. Programing and user interfaces are addressed in Section 4, and a demonstrative example of DISplay's use is given in Section 5. We conclude in Section 6, discussing the contributions of this work, aspects of its limitations and ideas for future work.




Figure 1: The Client Server Architecture

2 The DISplay Architecture

The DISplay software architecture is based on the client-server paradigm [2]. Here, the simulation application, which is created by an analyst or programmer, is treated as a client. Client calls are made by invoking functions resident in the DISplay client library. The server portion of the architecture consists of a **connection server (CS)** to which clients connect, and a **DISplay server (DS)** which handles all graphics and user interaction requests from a distributed application. Typically, the client connects to the **CS** during application initialization; the **CS** delivers a handle (i.e., socket) to the application for communication with the **DS**. The **DS** is created by the **CS** using a unix *fork* operation, and reads incoming messages from the application, creating the necessary display *tasks* and interaction *dialogs*, as specified by the application. Continuing display and user interaction is based on application specifics and objectives. Messages between the application and the **DS** form a well-defined set of primitives, part of the DISplay protocol. The **DS** connects to an X workstation to execute specified display and user interaction commands. The analyst can interact with an executing application from the workstation where the display is done, thus enabling a requisite interaction capability between analyst and model. A picture depicting the overall architecture can be seen in Figure 1.

The **CS** performs an important function in the mechanism described above. It must be located on a well known machine and port, so that applications can readily connect to it. The connection is made using a unique *Name*, which identifies the application. The **CS** maintains a list containing each active **DS**, along with its associated *Name*. In a multiprocessor application, say using PVM, Conch, or some other system, several processes, possibly residing on distinct (distributed) processors, may request a single, cooperative graphical display and user interaction. In this case, the **CS** delivers the same **DS** handle to each of these processes (from the same application), so that they may all connect to the same **DS**. All executing processes may be on distinct and remote machines, including both the **CS** and the **DS**, which need not be located on the workstation where the display and user interaction is being done.

Once a **DS** has been created, an application interacts only with its **DS**, leaving the **CS** to continue to

read connection requests from other applications. If the name specified in an incoming connection request is different from any of the current names in the **CS** database, it delivers the application a connection to a new **DS**. Otherwise, it delivers the application an existing connection. Such a scheme is the primary means for sharing a server among several cooperating processes, and is common in parallel computing environments.

2.1 Use of **DISplay** in Parallel Simulation Environments

There are several approaches to developing parallel discrete event simulations. Of these, the primary approaches are *conservative*, *optimistic* and *adaptive*. In the conservative approach events are processed strictly in order of occurrence, maintaining causality at all times. The approach is prone to deadlock, which may be avoided by resorting to the use of so-called *null messages*. In the optimistic approach, events are processed as they become available; potential causality conflicts are disregarded until a causality error is detected. Upon detection of such an error, invalidated simulation work is undone, the causality error corrected and the simulation re-executed from the point of correction. The approach requires some form of state-saving, so that simulation re-execution from a given point is possible. The adaptive approach is a mixture of the optimistic and the conservative approaches. A summary of these approaches can be found in [3].

Without loss of generality, we assume that the simulation environment is made up of a cluster of workstations which work collectively to solve a problem. Each workstation may host a variable number of processes. The **DS** can be used in one of two ways in such a situation. One approach is to connect every simulation process to the **DS**. Doing this, however, may not be possible because the **DS** may run out of file descriptors if the number of processes is very large, since there is a limit on the number of open descriptors at any given time. An alternate approach is for one or few of the processes to collect messages from other processes and pass these on to the **DS**. This approach suffers the additional overhead incurred by the intermediary or intermediaries, if such message collection is not an inherent part of the executing parallel application. For each process in a given simulation that connects to the **DS**, the latter creates a virtual channel on which messages from the sending process are queued. It is assumed that messages along each channel are received in order of simulation time.

2.1.1 Synchronization

During a parallel simulation, distinct processors may be involved in computations associated with different virtual simulation times. Messages sent to the **DS** by different processors arrive at the **DS** with distinct virtual-time stamps. To provide the user with a consistent view of the simulated system, the **DS** buffers, sequences and processes messages only when it is certain that subsequent messages cannot induce causality violations. One way of achieving this is by adopting a conservative parallel simulation protocol in **DISplay**. Consider an example utilizing the conservative protocol. Each channel (source-destination link between processes, as viewed from the destination end) is associated with a channel time. This time is equal to




Figure 2: Synchronizing DISplay with the application

the time-stamp of the first message queued on the channel buffer, if one exists, or the time-stamp of the last message retrieved from the channel buffer if the channel buffer is empty. At any given time, the first message from the buffer of the channel with the smallest channel time is selected for processing. The following description uses an example to simplify exposition.

In Figure 2 is shown a set of three processes, numbered 0,1, and 2, that connect to the same **DS**. The **DS** creates a virtual channel for each, using processor identifiers to distinguish between channels. Channel 0 has the smallest channel time, and so the message with time-stamp 25 in channel 0 is selected for processing. After this message is processed by the **DS**, further DISplay processing is temporarily halted because channel 0 becomes empty (with channel time still at 25), and a minimum time-stamp message cannot be identified. Now if, for example, another message arrives on channel 0 with a time-stamp of 27.5, channel time for channel 0 is updated to 27.5, and channel 2 wins as the channel with the smallest channel time. The message with time-stamp 27 on channel 2 is selected for processing, leaving channel 2 with a channel time of 28.5. As long as messages continue to arrive from processes, channel times can be updated, and a minimum time-stamp message selected for processing.

If a process has no messages to send, it avoids an indefinite delay of message selection at the **DS** by forwarding a Time Update message, which is basically the equivalent of a null message in the conservative approach. In effect, the processor sending the message informs the **DS** that no message will arrive before a time specified in the Time Update message. Observe that the **DS** may also be used in an optimistic parallel simulation, in which case messages will not be sent to the **DS** until they are committed and cannot be rolled back. This occurs when the time-stamp of a message is smaller than a parallel simulation's global virtual time. In this case however, Time Update messages are not required but synchronization is necessary so that a consistent view of the simulation is provided to the user. An

example of the use of interactive graphics in the optimistic simulation system SPEEDES is given in [20]. For applications that have no notion of time, it is feasible to send all messages with the same time-stamp of 0. However, requests must not be sent in wrong time order. If this occurs, there's a good chance that the application's execution logic is incorrect.

2.1.2 Resource Sharing

DISplay provides parallel applications with two basic resources, called *tasks* and *dialogs*. *Tasks* are used for displaying simulation output graphically, while *dialogs* enable user interaction. In a parallel execution environment, it is necessary for several processes to share a single resource at the server. DISplay permits *task* sharing between processes. For example, all computing processes may employ the same *task*, to perform displays in a single window. This is useful in domain-decomposition applications where processes work on different parts of a domain. Each process displays its contribution to an entire graphical result via a mapping to a relevant portion of the output screen. Such an approach can be useful when comparisons between processor results are required – distinct processes display results in the same window, and a simple visual scan of the display allows for easy comparison.

2.2 Design Goals

The DISplay software system (servers and client library) is part of a larger development effort geared towards parallel computations and simulations. This effort is manifest in the ACES software architecture, a brief overview of which can be found in [7]. A first motivation was the need for a general mechanism to display data and perform user interaction in a software layer that made display and interaction functions independent of kernel and communication layers of the ACES system. A second motivation for the DISplay tool was the need to perform graphical debugging and performance monitoring of the system's nontrivial function-interactions, both within and between layers. Facilities for debugging and monitoring are currently absent in DISplay, but we expect to include this functionality as the project matures.

While additional goals may be added as the system evolves, the basic goals of DISplay at present are: **Simplicity:** provide the user with a simple means to collect and view results from parallel simulations. This should relieve the user of the responsibility of making and maintaining connections and creating displays. **Extensibility:** provide an open library in that new functionality can be added with relative ease. **Interaction:** provide complete interaction, so that the application can send data to the user, and the user can send data to the application. This reciprocal action allows the application to take in input, give the user simulation output and also allows the user to change simulation variables dynamically, during a run. **Performance Monitoring:** provide a means to visually depict the status of a simulation at each process in a parallel simulation, so as to be able to identify critical sections or bottlenecks in the simulation. **Generality:** provide as general an interface as possible, so that the system can be used with general parallel applications. **Portability:** provide software that is minimally dependent on particulars of machines, that is inter-operable so that ports to a variety of platforms is possible.

3 The Display Server

The power of display server (**DS**) lies in its generality – because it does not store application-specific state. It is application-independent. The **DS** accepts information from an application for the purpose of display, and it is up to the user to make decisions on types of displays and types of interactions desired with an application. The **DS** presents information to the user in a well-defined manner, accepts input which it forwards to the application – if the application requires this input, and displays results – if the application requests such a display. All interpretations of presentation are left to the user. Functions that the **DS** provides execute independently of the application. For example, clicking on a button to display a graph will not affect the executing application. This function is handled at the **DS** itself. Since the **DS** is X-windows based, it provides the familiar windows, icons, menus, and pointing devices interface. It provides a point and click interface, for function invocation. These invocations may occur randomly, e.g. windows may be resized, uncovered, or iconified/deiconified, buttons that invoke user interactions or display windows may be clicked on.

The **DISplay** system provides *tasks* and *dialogs* as basic mechanisms for implementing graphical displays and user interactions, respectively. *Tasks* are of two types: Single message tasks require a single message to be delivered, and their action is taken based on that message alone. Multiple message tasks require a setup phase, where the task is created locally and then executed at the server. The creation mechanism returns a task identifier which is used as a handle in all subsequent messages related to the task. After task creation, multiple task related messages may be sent to the **DS** to carry out the intended function. These messages are interpreted in the context of the task identifier. Tasks are deleted by using an `EndTask` message. The use of a task identifier provides for some useful performance improvements. For example, it is not necessary to send information repeatedly saying that a plot is 2D or 3D, with every message that is associated with this plot. This information is stored at the **DS** at task creation time, and a request for plotting a point is executed based on whether the plot is setup for 2D or 3D.

User interactions are also handled in a similar manner. Each possible user-model interaction is assigned a unique dialog identifier, with a *form* based dialog created. A dialog setup phase creates the dialog locally and then sets it up at the **DS**. Queries and answers using these dialogs are interpreted based on the dialog identifier and the position of the data within a dialog. It is possible to link dialogs to specific tasks, and to have tasks linked to specific dialogs. For example, the node or arc of a graph task can be associated with dialogs. Also, when a dialog is activated it may simply display an active task like a plot of points.

The **DS** repeatedly reads messages (including requests for new connections from clients) on incoming channels and acts on them. Messages that are used for setup and termination are identified as `control` messages, and others are referred to as `action` messages. Each message has a header which contains the type of message, virtual time, task identifier, a process identifier for the sending process, the simulation identifier (which depends on name of the simulation, and is assigned by the **CS**), and an operation code. If the message is of a `control` type, action is taken immediately provided the operation code in the

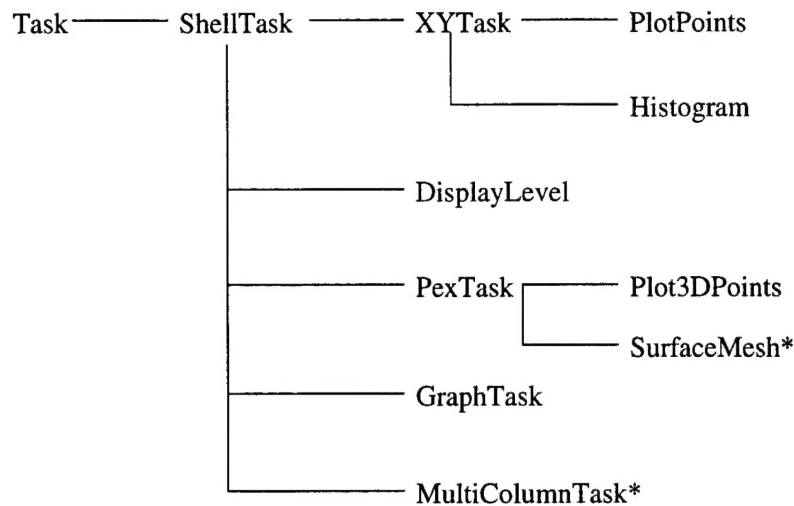


Figure 3: The C++ Task Hierarchy

message is *urgent* or specifies the creation of a new process. Otherwise, the message is queued and is acted on in correct time order. Queued messages are processed by a queued message handler which invokes the appropriate handling routine, based on message type and other relevant information contained in the message, e.g., task identifier, dialog identifier. Commands to display output or query the user are then executed. Details on these aspects of the system can be found in [8].

3.1 DISPLAY Classes

A central class, called the *controller* class, keeps track of all connected processes and their associated channels. It has a predefined maximum number of channels which are setup as each process connects to the **DS**. Each *channel* class records the socket over which its communication takes place, channel virtual-time, and a list of tasks and dialogs indexed by identifiers. Messages received on the channel are queued in FIFO mode. Other important classes in the system include the *task* and *dialog* classes.

3.1.1 Tasks

DISPLAY Tasks are used for any output requiring multiple messages. *Tasks* form an extensible set of textual, graphical and representational schemes for displaying the dynamics of a simulation and its end results. The C++ class hierarchy shown in Figure 3 describes the arrangement of tasks. The *Task* class provides control functions for maintaining the status of a task. A task may need to be displayed even after the simulation is complete. Thus, the display does not disappear when the simulation terminates, and may correspond to the final simulation result. When an *EndTask* control message is received from a client, the task is marked for deletion. It is deleted only when the user selects the *Cancel* button from the task window on the screen. When new task classes are added, these may utilize the functionality

Interaction	Dialog Invocation	
	<i>Model Prompted</i> (Synchronous)	<i>User Prompted</i> (Asynchronous)
<i>Two Way</i>	Model Query Dialogs	User Query Dialogs
<i>One Way</i>	Model Display Dialogs	User Command Dialogs

Table 1: Classification of DISplay Dialogs

provided by the base task classes, such as the Task class. This facilitates extension of the interface with newer displays. The ShellTask class provides for the popping up and down of the window in which the display is drawn. XYTask is used for all 2D tasks that are drawn to scale. PexTask contains basic data structures for use with tasks that use the PEXlib 3D library. At the leaves are the actual classes that a user must be aware of, though it is never necessary to program directly with these.

The classes marked * in figure 3 are to be added in the near future. For example, a provision for the display of multicolumn tables is planned. This may easily be sub-classed from the ShellTask class. Adding a new task means that the DISplay protocol must be expanded to include the new task message. Hooks must be provided in existing DS functions so that control may be passed to functions which will act upon these messages. A new set of functions pertaining to the new functionality must also be added to the client library so that the simulation analyst can make use of the new task.

3.1.2 Dialogs

DISplay *dialogs* are basic mechanisms for user interaction. A *dialog* is defined as a *form* consisting of a set of values. Associated with each value is a set of items which describe that value. In the current implementation, this set of items is limited to a textual description of the item and a specification describing the type of the value. The textual description can also be used as a query to the user to input a value of the required type. Other qualifiers may be added to the existing set. For example, in order to do validation of user input, a valid range for the values may be provided. Several related dialogs may be made part of a single *panel* from which the user can point and click on the label of the particular dialog that is being invoked. On being invoked, the current state of the dialog is displayed. Dialogs allow for bi-directional interaction between the application and the user.

Dialogs are of four types (see Table 1), depending on the direction in which interaction is allowed or restricted, and whether the dialog is invoked by the model or the user.

1. **User Query Dialog.** The request for value information is sent from the DS to the application. The application handles the query and returns the required values. Invocation is done by the user, who may fill in some of the values in the dialog. The application may use these in computation, returning computed values.

2. **User Command Dialog.** The user sends information to the application asynchronously. The application must be willing to handle these dialog messages through appropriate user written handlers which are registered with the dialog on creation.
3. **Model Query Dialog.** This functionality is similar to remote calls. The application pops up the dialog at the **DS** at predefined points in its execution, and waits for the user to reply. The application remains blocked until the user reply is received.
4. **Model Display Dialog.** Information is passed to the **DS** at predefined points in the application execution. The **DS** collects the information and displays it only when the user invokes the dialog (by clicking on the associated button). The information is received only when the application chooses to send. A special form of this dialog is one associated with a task. The associated task is popped up when the panel button is clicked.

It is the developer's responsibility to provide for functions called *dialog handlers* in the application. These must handle replies or queries from the user and are registered when the dialog is created. They are called automatically when a dialog message of the given type is received. If asynchronous dialogs are used, the developer must arrange to call a library-provided function from time to time, to read dialog messages from the **DS** and call application dialog handlers. An example given in Section 5 shows how this may be done. Since all messages are logged, it is possible to do a post-run determination of times of interaction between user and application. This may be important in performance-tuning and analysis of simulations.

4 Programming and User Interfaces

4.1 The Programming Interface

There are two simple approaches to using **DISplay** with an application. One approach is to first construct the simulation application, without regard to **DISplay** interaction. Calls to the **DS** can be added when the working simulation is available. With this approach, the user cannot avail of development support and debugging help (i.e., display of run-time status, values of simulation variables and simulation object states) provided by **DISplay**. On the other hand, if parallel debuggers are available, they may compensate for or even offer services superior to **DISplay**'s debugging services. Good parallel debuggers, however, are still years away. Given a working simulation, client calls may be added wherever necessary in the application code.

An alternate approach, which we recommend, is to create the simulation application *with* provision for client calls in the application design. In testing simulation logic, the **DS** may be disabled – either by not initiating a connection to the **CS** or by setting the value of the handle returned by the `sGetServer()` call to **SDERROR**. This constant is defined in an include file **sdconsts.h**, which must be included along with header file **sdclnt.h** in all programs that make calls to the **DS**. With the handle thus set, send-data

calls to the server return immediately, with no effect on the simulation. The first example described in the following section was developed using the first approach simply because working simulation code was already available. The second example described was developed using the second approach; the **DS** was disabled until the application was tested.

To utilize **DS** functions the application developer must invoke functions in the DISplay Client Library (DCL) `libds.a`, which must be linked with the application. The DCL consists of C functions – so it can be used with any application that can be interfaced with the C language. All DISplay functions return appropriate error indicators, and functions that fail return `SDERROR`. For example, the availability of the **DS** is determined by examining the value returned by `sGetServer()`. Subsequent calls to the DCL may first ascertain the availability of the **DS** by testing this value. This allows the simulation to execute even when the **DS** is unavailable. If interactive use of the simulation is to be replaced by batch use, the analyst simply needs to set the window parameter in setup calls to `FALSE`. Of course, this will work only if synchronous dialogs (which require user replies) are not built into the application, or are replaced by calls to read from a file depending on the value of a switch. Logging is also allowed in batch-mode, with messages to the **DS** logged in DISplay message format, to be replayed at a later time using a DISplay utility program.

Client functions are of two types: functions which send messages to the **DS** and functions which perform local processing prior to sending a message. Functions of the former type begin with an `s` prefix, while functions of the latter type begin with an `sd` prefix. The DCL contains a host of basic functions, with several convenience functions that invoke basic functions. An application developer may call either basic or convenience functions. The latter are simpler to use and recommended, whenever possible. At present, the basic function set consists of the 34 functions listed in Table 2 shown in the Appendix. These are classified according to functionality. The first two arguments to all message-sending functions consist of a socket identifier and simulation time. The third argument is usually a task identifier or a dialog identifier. Figure 7 in the Appendix gives an example showing the typical use of functions to connect to the **DS**. In setting up the connection, it is possible to specify whether logging is required, and whether a new main window is required. By specifying the number of processes in the simulation, all processes are guaranteed connection before the **DS** begins to act on messages from the simulation.

Once the connection is made, tasks may be created, initiated and terminated at any time. In Figure 8 (in the Appendix) is shown the use of functions for creating a Histogram task (`histtask`) and two dialogs. The first dialog (`dialog_task`) is associated with the newly created task `histtask`. The second dialog (`dialog_sync`) is a synchronous dialog which is invoked using the `sQueryReply()` call. The global argument to function `sdCreateTask()` ensures that different processes will cooperate to draw in the same window. For example, in a particle dynamics application the domain is set up as a 2D grid, divided into equal-sized slices among processes. Though processors compute on distinct domains, it is possible to display the particle movement on the entire grid in a single window. All processes write into a globally known window, each writing onto only that portion of the window for which it is responsible, which may indeed be the entire window. This is preferable to the situation where processes

create and display particle movement in distinct windows – this does not aid result interpretation and analysis, especially when the number of windows is large.

The functions which do the brunt of the drawing work in task windows are listed as *Multi Message Task Functions* and *Graph Message Functions* in Table 2. These functions take in a task identifier as a parameter, to identify the task to which they must refer on the **DS** side. Important functions for drawing in a window include `sColPoint()`, `sPltPoint()`, `sHistPoint()`, and `sDisplayLevel()` – capable of 2D and 3D drawings. Function `sColPoint()` colors a point in the window with a specified color. The programmer programs with world coordinates and color names. The **DS** converts world coordinates into corresponding window coordinates, and maps colors to pixel values that can be displayed on the specified X workstation. Function `sPltPoint()` plots a line between two points. The corresponding convenience function for performing 2D plots is `sPltPoint2D()`. Function `sHistPoint()` places a histogram line on the task window, and function `sDisplayLevel()` displays the specified level of a meter or variable. With the aid of such functions, diverse abstract displays like scatter plots, line plots, histograms, x-y-z plots, and level indicators may be shown. Some of these displays are *instantaneous* displays (the display shows the status of the simulation at specific points of time), whereas other displays are *cumulative*, in that users are presented with a history of information (for example, a x-y plot task may display a continuous change of the value of a simulation variable plotted against time).

DISplay allows representational displays to be created using networks. Networks can be used to represent displays for a variety of physical systems. In queueing systems graph nodes may represent servers, with arcs representing possible customer routing between servers. In most cases some form of visual representation of the physical object is used for a graph node. The default is to use a rectangular box with a label. It is possible to associate dialogs with nodes and arcs. Clicking on a node or arc with which a dialog is associated, will cause the associated dialog to execute at the **DS**. For example, in an *EcliPSe* simulation, processors arrange themselves in a virtual tree-topology [9] to perform parallel sampling. This underlying tree configuration of the machines can be displayed by DISplay's Graph Task (see Figure 4). With each node is associated a dialog that displays performance data for that node, causing it to pop up when that particular node is clicked. Moreover, the graph can be modified dynamically. Nodes and arcs may be added or deleted. The colors of the nodes and arcs may be changed to signify a change in the value associated with the node (e.g., CPU load level) or arc (e.g., amount of traffic).

4.2 The User Interface

The **DS** provides the analyst with a user interface that can be tailored to suit the application. This graphical interface is common to all DISplay applications. For example, a sample screen dump of an *EcliPSe* [9, 7] Performance Monitoring display is shown in Figure 4. This is a particular use of DISplay for *EcliPSe* applications. In the figure is shown an example of a network (GRAPHTASK), a histogram task and several asynchronous dialogs. The interface consists of the *Interaction Push-buttons* which, when clicked, cause

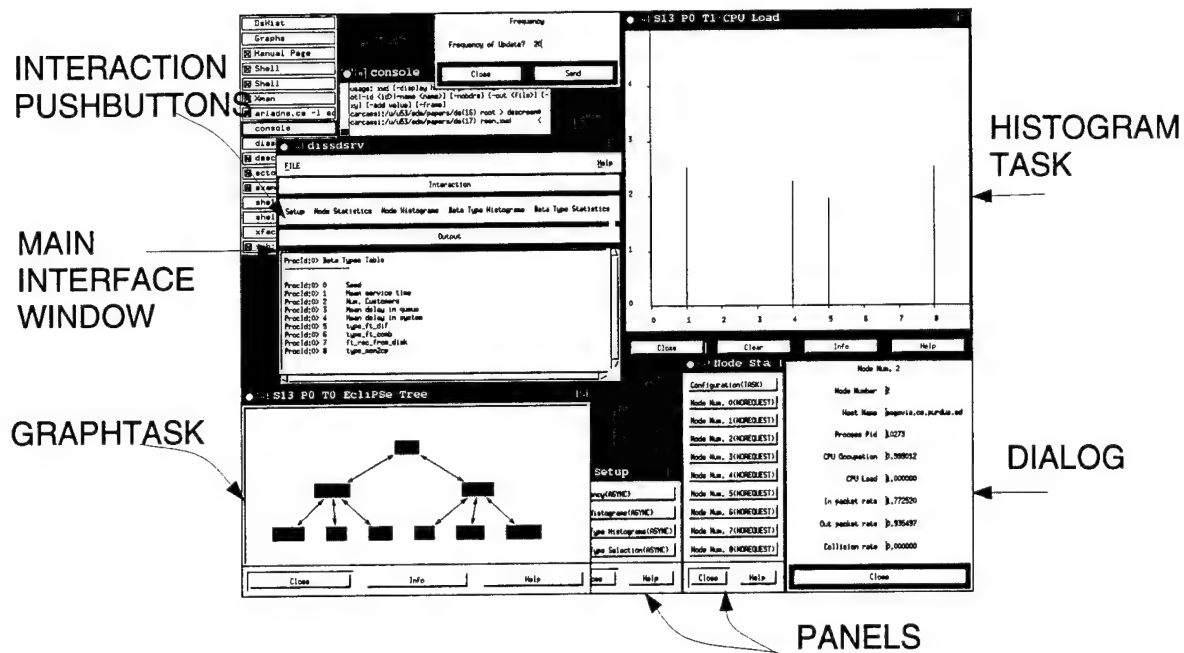


Figure 4: The DisplA User Interface

Panels to pop up. Each panel encapsulates a group of associated dialogs or displays which may be invoked through use of the panel buttons. Button labels clearly indicate button functions. Each interaction push-button displayed on the main menu bar of the user interface is application specific. To modify a user-presentation (for example, to add another panel option) it is necessary to incorporate the option in the program code and recompile the user's application. Thus creating application specific interfaces does not require the client library or the server software to be modified. Task windows open up whenever a Task is initiated in the application. The X window system allows the user to arrange multiple windows on the screen, as required. The **DS** also handles redraw and resize commands independently of the application. Scale sizes and limits may also be changed interactively and the complete window is redrawn with the new sizes and limits. Windows can be iconified and deiconified as necessary, and dialogs can be invoked to display data and closed when not required.

5 Example of DISplay Use

The DISplay system software is useful in a variety of situations including product demonstration, gaming, learning, modeling, performance monitoring, parallel debugging, etc. The use of the **DS** for monitoring of parallel applications in *EclIPSe* is described in [7]. Here we focus on a simple example demonstrating the use of DISplay in general simulations. The example is that of an M/M/n queueing system. This example shows how user interaction facilities and basic displays are used. Results are displayed using line plots and histograms. It is instructive to note that the user can dynamically alter simulation variables such as

customer arrival rate and/or number of servers, to examine system behavior under such changes.

5.1 M/M/n Queueing System Simulation

An M/M/n queue is an n -server queue with Markovian arrivals and services. Customers queue in FIFO mode in a single queue, awaiting service. In this example we fix the mean service time μ and examine the effect of arrival rate λ and number of servers n on the measured responses of mean queue size, actual queue size and server utilization. The M/M/n application was developed using the process-oriented simulation language CSIM [16]. The multi-server facility is implemented using the `facility` type in CSIM. This simulation tool also provides functions for computing queue size, mean queue size and utilization, all of which were used to collect data that was subsequently sent to the **DS**.

In this example application, a user interface was rapidly generated using three dialogs. One was a synchronous, model-prompted dialog, where the application blocks and queries the user to enter input parameters (i.e., λ , μ , n and total number of customers to be simulated). Once these parameters are initialized, the application continues to execute. The other two dialogs are asynchronous, user-command dialogs. These allow the user to dynamically alter λ and n . New values for these parameters are sent to the application, which echoes back both new and old values to the user interface. Appropriate handlers must be provided in the application to act on these parameter changes.

To display the state of the system three displays are used. The first display depicts the utilization level of the multi-server facility. The second display shows both the mean queue size and the actual queue size on a single graph. The third display shows a queue size histogram, made by updating samples on each customer departure from the facility.

The different dialogs just described are shown in Figure 5, with the synchronous dialog displayed in Figure 5(a). The code segment shown in Figure 9 (in the Appendix) shows how this dialog is set up and used. Each entry in the dialog form has a corresponding entry in the `SdQtnAns` structure. The dialog is set up locally using `sdCreateDialog()`, with `SdQtnAns` passed to the latter as a parameter. Dialog initiation at the **DS** is accomplished with `sBeginDialog()`, and invocation via the function `sQueryReply()`.

Asynchronous dialogs are set up locally and initiated at the **DS** in the same manner. But in contrast to synchronous dialogs, these are initiated when the user clicks on specific buttons that appear in the user interface window. Appropriate handlers and calls to read messages from the **DS** must be provided. In the application, the developer must call the DCL function `sHandleReply()` periodically, to look for queries or commands from the user, with appropriate handlers invoked in response. An example of a handler (`handle_arrival_change()`) whose function is to change mean inter-arrival time on user input is shown in Figure 9. This example shows how application-specific dialogs are constructed and used; dialogs can also be changed with ease, if such a change is required. Dialog handling code is easily separated from the main part of the application through use of handlers that are automatically called when a dialog is invoked by the user.

M/M/n Parameters	
Mean Inter-Arrival Time	2
Mean Service Time	4
Number of Servers	5
Number of Customers	80000
This is a query from an application - you must reply!	
<input type="button" value="Close"/> <input type="button" value="Send"/>	

(a)

Number Of Servers	
Change Number of Servers to:	
Old Number of Servers was:	
<input type="button" value="Close"/> <input type="button" value="Send"/>	

(b)

Mean Arrival	
Change Inter-Arrival to:	
Old Arrival Rate was:	
<input type="button" value="Close"/> <input type="button" value="Send"/>	

(c)

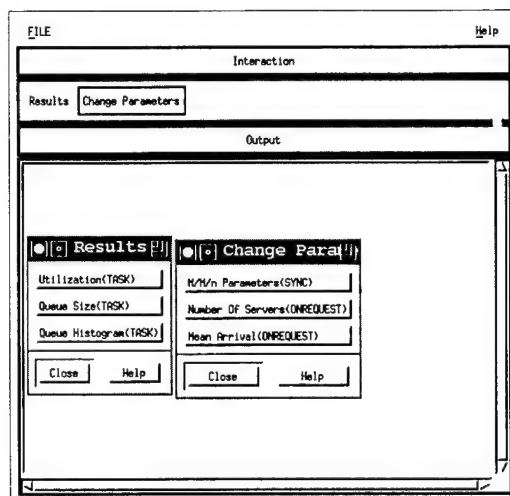
Figure 5: Dialogs for the M/M/n Server System

Figure 6 contains a user interface window for the queueing application, and also windows which display simulation results dynamically during model execution. In Figure 6(a) is shown the main user interface window, with two interface push-buttons corresponding to Results and Change Parameters. By clicking on these buttons, the user obtains panes with dialog-activating push-buttons. These panes are shown in the Output portion of the window in Figure 6. Results are displayed in abstract form, often with plots and histograms. For example, in the queueing application, model parameters are initially set as depicted in the dialog given in Figure 5(a). In Figure 6(b) can be seen a simple plot of facility utilization versus time. In Figure 6(c) is shown a realization of queue size as well as mean queue size, with both graphed against time. Figure 6(d) shows a dynamically computed histogram of the queue size.


The graphs show that utilization can be increased, during a simulation run, by interactively changing the arrival rate λ ; this also results in a larger mean queue size. Facility utilization is seen to fall to zero at certain times. These times coincide with the times at which the number of servers in the CSIM facility was changed. For example the points labelled A in Figure 6(c) correspond to points where λ was changed, and the points labelled N correspond to points at which the number of servers was changed. Subscripts on each symbol give the new value, after the change.

6 Conclusion


The DISplay software is an application-independent tool for performing Visual Interactive Simulations. It may be used in parallel simulations or computations. It provides the model developer with tools for parametric description of a large set of graphical outputs and capabilities for data visualization, inspection




(a) Main Interface Window



(b) facility Utilization



(c) Queue Size and Mean Queue Size(MQS)



(d) Queue Size Histogram

Figure 6: User Interface Window and Results

and specification. Abstract as well as representational displays are implemented and these can be dynamic or static in nature. Provision is also made to capture the instantaneous state or the cumulative state of an application. User interaction allows for intervention in model execution, with two way interaction permissible. Model specific handlers can be invoked automatically when a particular type of interaction is invoked by the user. Specific support for parallel computations is provided by allowing several processes to cooperate and draw in a common shared window, and to synchronize with the display. DISplay provides the model developer with a flexible yet powerful tool to create application specific displays and user interactions without any need to know the underlying complexities of graphics and data communications in a distributed environment.

The system described can be enhanced in several ways. Use of a specification file for task and dialog descriptions would be useful in setting up these resources, instead of coding these directly into the program. Interface modification would thus only entail change in the specification file. The specification file could be built interactively. Several tasks, including support for three dimensional surface meshes, and display of multicolumn tables can be added; display of results at more than one workstation — to provide distributed displays — so that multiple users can collaborate in simulation analysis also would undoubtedly be useful. These enhancements are currently under way, as part of a larger effort in heterogeneous distributed simulation. This effort is based on the ACES system [7].

The basic types of simulation output described here are common to most simulations. This realization motivated us to develop a generalized and extensible display capability. DISplay is now a full fledged VIS tool, enabling easy user interaction and dynamic visualization of output in nontrivial applications. We have used DISplay to visualize parallel simulated annealing, numerical analysis algorithms, particle-physics, and computer network simulations.

References

- [1] N. Collins and C. Watson. Introduction to Arena. In G. Evans, M. Mollaghasemi, E. Russell, and W. Biles, editors, *Winter Simulation Conference*, pages 205–212, December 1993.
- [2] D. Comer and D. L. Stevens. *Internetworking with TCP/IP*, volume III: Client–Server Programming and Applications. Englewood Cliffs, N.J. : Prentice Hall, 1993.
- [3] R. Fujimoto. Parallel Discrete Event Simulation. *CACM*, 33(10):30–53, 1990.
- [4] T. Gaskins. *PEXlib Programming Manual*. The Definitive Guides to the X Window System. O'Reilly & Associates, Inc., 1992.
- [5] M. Glavach and D. Sturrock. Introduction to SIMAN/Cinema. In G. Evans, M. Mollaghasemi, E. Russell, and W. Biles, editors, *Winter Simulation Conference*, pages 190–192, December 1993.

- [6] D. Heller. *Motif Programming Manual. For OSF/Motif Version 1.1*, volume Six of *The Definitive Guides to the X Window System*. O'Reilley & Associates, Inc., motif edition, 1991.
- [7] F. Knop, E. Mascarenhas, V. Rego, and V. Sunderam. Fail-Safe Concurrent Simulation with EclIPSe: An introduction. *Simulation Practice & Theory (to appear)*, 1995.
- [8] E. Mascarenhas and V. Rego. DISplay: A technical reference manual. Technical report, Purdue University, Department of Computer Sciences, 1994. In preparation.
- [9] H. Nakanishi, V. Rego, and V. S. Sunderam. On the effectiveness of superconcurrent computations on heterogeneous networks. *Parallel and Distributed Computing (to appear)*, 1994.
- [10] R. M. O'Keefe. What is Visual Interactive Simualtion? (And is there a methodology for doing it right?). In *Proceedings of the Winter Simulation Conference*, pages 461–464, 1987.
- [11] J. J. O'Reilly. Introduction to SLAM II and SLAMSYSTEM. In G. Evans, M. Mollaghasemi, E. Russell, and W. Biles, editors, *Winter Simulation Conference*, pages 179–183, December 1993.
- [12] J. K. Ousterhout. An X11 Toolkit based on the Tcl language. In *Winter USENIX Conference*, pages 105–115, 1991.
- [13] Pritsker and Associates, Inc., West lafayette, IN 47906. *TESS and SLAMII*.
- [14] M. Rooks. A Unified Framework for Visual Interactive Simulation. In *Proceedings of the Winter Simulation Conference*, pages 1146–1154, 1991.
- [15] E. C. Russell. Building simulation models with SIMSCRIPT II.5. CACI Products Company, La Jolla, CA, 1989.
- [16] H. Schwetman. *CSIM Users' Guide*. Microelectronics and Computer Technology Corporation, June 1992.
- [17] V. Sunderam. PVM: A framework for parallel distributed computing. *Concurrency: Practice and Experience*, 2(4), December 1990.
- [18] W. Thompson. A tutorial for modelling with the WITNESS visual interactive simulator. In G. Evans, M. Mollaghasemi, E. Russell, and W. Biles, editors, *Winter Simulation Conference*, pages 228–232, December 1993.
- [19] B. Topol. Conch: Second generation heterogeneous computing. Master's thesis, Department of Math and Computer Science, Emory University, 1992.
- [20] Y.-W. Tung and J. Steinman. Interactive graphics for the parallel and distributed computing simulation. *Proceedings of the Winter Simulation Conference*, pages 695–700, 1992.

A Appendix

```
main(int argc, char **argv)
{
    char *simname;           /* the name of the simulation */
    char *host;              /* the location of the connection server */
    char *service;          /* the port number of the connection server */
    int sock;               /* the handle to the DS */
    int ret;                /* return code from functions */
    int window=TRUE, log=TRUE; /* window up and logging on */
    int num_of_procs;       /* number of processes in the simulation */

    ...;

    sock = sGetServer(simname, host, service); /* make the connection */
    if (sock == SDERROR) {
        /* handle the error here */
    }
    /* register this process as one of num_of_procs */
    ret = sNewProcessMsg(sock, window, log, num_of_procs);

    /* rest of processing and messages to DS */

    ret = sEndServer(sock);      /* close the connection */
}
```

Figure 7: Example of connecting to the Display Server

```

#include "sdconsts.h"
#include "sdclnt.h"
#include "sdutil.h"
{
    /* dialog */
    static SdQtnAns qtnans_sync[] = {
        {"What is your name?", SDSTRING},
        {"What is your key?", SDINT},
    };
    char *reply[2];
    double timereply=0.0;
    int i;

    int histtask;
    int dialog_task, dialog_sync;
    int global = 1;          /* is a global task */

    /* create and initiate a Histogram task */
    histtask = sdCreateTask(HISTTASK, window, log, global);
    sdSetTaskValues(histtask,
        SdXaxisName, "X- axis",
        SdYaxisName, "Y- axis",
        SdTitle, "Particle simulation",
        SdXmin, 0.0,
        SdXmax, 20.0,
        SdYmin, 0.0,
        SdYmax, 20.0,
        SdNoPoints, 5.0,
        SdDim, 2,
        SdXaxisInterval, 1,
        NULL);
    sBeginTask(sock, simtime, histtask);

    /* create the task dialog */
    dialog_task = sdCreateDialogTask("Histogram", histtask);
    /* create the Synchronous dialog */
    dialog_sync = sdCreateDialogNotask("PassKey", SDMODEL_QUERY,
        Number(qtnans_sync), qtnans_sync,
        NULL);

    /* begin the dialogs grouping the two into a single pane*/
    sBeginDialog(sock, simtime, dialog_task, dialog_sync, "Some Dialogs");

    /* invoke the synchronous dialog */
    ret = sQueryReply(sock, &timereply, dialog_sync, reply);

    /* print the replies */
    for (i = 0; i < Number(qtnans_sync); i++) {
        printf("Qtn:%s Ans:%s\n", qtnans_sync[i].question,
            reply[i]);
    }
}
}

```

Figure 8: Use of Task and Dialog Functions

```

int sock;                                /* socket for communication */
int iatm;                                /* the mean inter-arrival time */
static SdQtnAns dialog_parm[] = {        /* parameters of the simulation */
    {"Mean Inter-Arrival Time", SDFLOAT},
    {"Mean Service Time", SDFLOAT},
    {"Number of Servers", SDINT},
    {"Number of Customers", SDINT},
};

static int parm_dialog_id, arr_dialog_id; /* dialog identifiers */

int ds_setup()                           /* demonstrates setup */
{
    ...;
    parm_dialog_id = sdCreateDialogNotask("M/M/n Parameters", SDMODEL_QUERY,
                                          Number(dialog_parm), dialog_parm,
                                          NULL);
    ...;
    sBeginDialog(sock, clock, parm_dialog_id, arr_dialog_id,
                 "Change Parameters");
    return(sock);
}

get_parameters(int sock)                  /* demonstrates use of Synchronous dialog */
{
    char *reply[4];                       /* replies placed here */
    double dbl, timereply;
    int ivl;

    timereply = clock;
    sQueryReply(sock, &timereply, parm_dialog_id, reply); /* query user */
    if ((dbl = atof(reply[0])) > 0.0)
        iatm = dbl;                        /* Inter-arrival Time */
    if ((dbl = atof(reply[1])) > 0.0)
        svtm = dbl;                       /* service time */
    if ((ivl = atoi(reply[2])) > 0)
        ns = ivl;                         /* number of servers */
    if ((ivl = atoi(reply[3])) > 0)
        nars = ivl;                      /* number of customers */
    sStrPrintMsg(sock, clock, "M/M/n parameters");
    sStrPrintMsg(sock, clock, "Service Time %f Inter-Arrival %f Servers %d",
                 svtm, iatm, ns);
}

int handle_arrival_change(char *reply[]) /* handle mean inter-arrival */
{                                       /* reply contains user input */
    float new_iatm;
    char **send_reply;

    if ((new_iatm = (float)atof(reply[0])) == 0) || (new_iatm == iatm)
        return(0);
    wait(event_list_empty);
    sStrPrintMsg(sock, clock,
                 "Changing inter-arrival time at time %f\n", clock);
    send_reply = sdCreateReply(arr_dialog_id); /* create a reply */
    sdAddToReply(arr_dialog_id, 0, new_iatm); /* place reply here */
    sdAddToReply(arr_dialog_id, 1, iatm);
    iatm = new_iatm;
    sReplyDispMsg(sock, clock, arr_dialog_id, send_reply); /* update screen */
    return(1);
}

```

Figure 9: Handler and Code for Dialogs in a CSIM program

<i>Function</i>	<i>Remote</i>	<i>Description</i>
Server Control Functions		
sGetServer()	Yes	Obtain a handle to the DS
sNewProcessMsg()	Yes	Identify itself as a process and request window setup and logging, if required.
sEndServer()	Yes	Close the connection to the DS
Single message Task Functions		
sNullMsg()	Yes	Send an empty message and update channel time
sStrPrintMsg()	Yes	Format and print the message on remote window
Multi message Task Control Functions		
sdCreateTask()	No	Set up a task and return a handle
sdSetTaskValues()	No	Set up specific parameters for the task
sdGetTaskType()	No	Get the type of task
sBeginTask()	Yes	Set up the task at the DS
sEndTask()	Yes	End the task at the DS
Multi message Task Functions		
sdCreateSdPoints()	No	Create a set of points to be filled in
sdFreeSdPoints()	No	Free the set of points created by sdCreateSdPoints()
sdAddSdPoint()	No	Add a point to the set created by sdCreateSdPoints()
sdFreqHistPoint()	No	Return the histogram frequency for a given X value
sColPoint()	Yes	Color a point with a particular color
sPltPoint()	Yes	Plot a line from the previous point to this point in a particular color
sHistPoint()	Yes	Draw the line bar in the Histogram
sRestartPlot()	Yes	Begin drawing a line plot from a different point
sDisplayLevel()	Yes	Display the level in a level-plot
Graph message Functions		
sGAddNode()	Yes	Add a node to the graph
sGAddArc()	Yes	Add an arc to the graph
sGDeleteNode()	Yes	Delete a existing node from the graph
sGDeleteArc()	Yes	Delete an existing arc from the graph
sGChangeVal()	Yes	Change some display parameter of a node or arc
Dialog Control Functions		
sdCreateDialog()	No	Create a dialog
sdDeleteDialog()	No	Delete a dialog when it is no longer required
sBeginDialog()	Yes	Set up a created dialog at the DS
sChangeDialog()	Yes	Change the parameters of the dialog
Dialog Handler Functions		
sdCreateReply()	No	Prepare a reply to send to the DS
sdAddToReply()	No	Add an answer in the reply created
sdFreeReply()	No	Free reply created by sdCreateReply()
sQueryReply()	Yes	Send a query to the User and get reply
sHandleReply()	Yes	Read user command and handle it in the application
sReplyDispMsg()	Yes	Send a reply to a User Query message

Table 2: Basic Functions used with the Display Server